Alternative Parallel Ring Protocols*

R. Mukkamala E.C. Foudriat K.J. Maly V. Kale
Department of Computer Science
Old Dominion University Norfolk, Virginia 23529-0162

April 17, 1990

Abstract

Communication protocols are known to influence the utilization and performance of communication network. In this paper, we investigate the effect of two token-ring protocols on a gigabit network with multiple ring structure. In the first protocol, a node sends at most one message on receiving a token. In the second protocol, a node sends all the waiting messages when a token is received. The behavior of these protocols is shown to be highly dependent on the number of rings as well as the the load in the network.

1 Introduction

The use of parallel communication channels to achieve a gigabit network is a very interesting concept. Especially, if a gigabit network needs to be built from an existing common carrier system, a network of parallel channels may be a viable alternate. However, its appropriateness can only be ascertained after determining its behavior under different load conditions. To this end, we chose a parallel ring network operating on token-based protocols.

Currently, our studies are restricted to two token-ring protocols. With each of the protocols, a token is assigned to each of the rings in the network. All tokens rotate in the same direction. While a node is holding a token for transmission, it cannot hold any other token.

1. Exhaustive Policy: Under this policy, when a node obtains a token, it transmits all the messages in its queue, and then releases the token.

*This work was supported in part by CIT under grant INF-89-002-01, by NASA under grant NAG-1-908, and by Sun Microsystems under RF596043.

2. Non-exhaustive Policy: Under this policy, a node can transmit at most one message when it receives a token.

In this report, we describe the results obtained from simulations of these two protocols. From these results, we make some comments regarding the appropriateness of parallel ring structures operating at gigabit speeds.

2 Input Parameters

In order to determine the behavior of the multiring token-based protocols, we have run a number of experiments. Following is a summary of the input data:

- Since we are only interested in gigabit networks, we have considered the total bandwidth of the network to be 10^9 bits/sec.
- When there are R rings in the system, each ring has a bandwidth of $10^9/R$ bits/sec.
- Since token-ring protocols are only relevant for local area nets, we considered a total ring length of 30 Km, with 30 equally spaced nodes.
- The propagation delay on the network is taken to be half the speed of light (i.e. 150Km/msec).
- For simplicity, we assume constant length messages (10K bits). Each message is transmitted as one entity (i.e. no fragmentation).
- The load on the network is expressed in terms of mean-time between arrivals (m) of messages at any node. Low loads are represented by $m = 1.0$, which denotes that one the average a message may be expected once every 1 msec at each node. High load is represented by $m = 0.31$. We have used $m = 0.5, 0.4, 0.35$ as other values. Message arrivals are assumed to be Poisson.

3 Results

The performance of the protocols is based on simulation of the network. The simulation was carried out for 10 seconds. The system was let to stabilize in the first 5 seconds, and the statistics were then taken in the second half of the experiment.

We make the following observations from the obtained results:

- In Figure 1, we compare the average time that a message at the *head* of a queue at a node needs to wait before a token on any of the rings is captured. This time is also generally referred to as residual time of the token inter-arrival time (RI_T). When the inter-arrival time of messages (at each node) is at least 0.5, the residual time is less than 0.3 msec. As the mean-time between arrivals decreases, this time increases significantly. Certainly, having multiple rings results in a reduced RI_T.

- Figure 2 illustrates the effect of number of rings on the average waiting time of a message. Under the exhaustive policy, the average waiting time (RT_T) is less than 1 msec for $m \leq 0.35$. The benefit due to the presence of multiple rings is more apparent at high loads. For low loads ($m \geq 0.5$), the average waiting time is almost the same as RI_T. This is obvious since the queue sizes are generally very small.

  Under the non-exhaustive policy, single-ring networks cannot tolerate situations where $m \leq 0.5$. (This can also be proved analytically). When there are at least 8 rings, the system can tolerate values of $m \geq 0.35$. When $m = 0.31$, however, the system has very large queues (hence not shown in the figure), resulting in very large average waiting times.

- Figure 3 summarizes the relationship between average token interarrival times and the number of rings. Obviously, the average token inter-arrival time ($E[T]$) at a node decreases with the number of rings, since the number of tokens is now increased. This decrease is more pronounced at high loads ($m \leq 0.31$). Under the non-exhaustive policy, the token inter-arrival times are quite low (even under high loads). This is not surprising knowing that in this protocol at most one packet per node per token is only transmitted. With the exhaustive policy, however, the token inter-arrival times for the tokens could be significant for smaller number of rings.

- Figure 4 illustrates the randomness of the measured inter-arrival times of tokens. This randomness is expressed in terms of the ratio of the standard deviation to the mean of the token inter-arrival times. The ratios are higher for the non-exhaustive policy. Since the non-exhaustive policy sends at most one packet at a time, while the exhaustive policy sends all the pending packets with a token, this observation is counter intuitive. We are attempting to explain this phenomenon through some probabilistic analysis. Generally, this ratio seems to increase with the number of rings (at least up to 8 rings). Beyond eight rings, the behavior of this ratio seems to depend on the load factor (m).

- Figure 5 describes the relationship between response time and the number of rings for different loads. Response time includes the time to wait in the queue, the time for transmission, and the time for propagation from source node to destination node. For low loads, since the waiting times are approximately constant, there is a slight increase in the response time with the increase in number of rings. The increase in response time may be explained by the increase in transmission delay due to reduced bandwidth per ring. The behavior of the average response time for both non-exhaustive and exhaustive policies is similar to that of the average waiting times in the queue.

- Figure 6 summarizes the variances in the response time. This seems to be quite different from the variances in the waiting times. First, the variances of the exhaustive and non-exhaustive policy are now comparable. Second, the variances are much less than those in Figure 4. Once again, we are investigating the possible causes for this phenomenon.

- Average token rotation time for each ring is also an important performance metric. This metric is shown in Figure 7. In the case of exhaustive policy, as expected, the token rotation time is independent of the number of rings (since increase in number of rings also increases the transmission delay by the same extent). The behavior of the non-exhaustive policy needs more investigations.

- Determining the probability with which an arriving token is used to transmit messages (f_T) is useful in describing the utilization of the network. This relationship is summarized in Figure 8. Clearly, the behavior of f_T differs under the exhaustive and non-exhaustive policies. For low loads (e.g. $m = 1.0$), both policies exhibit similar behavior.

- Figure 9 summarizes the behavior of the average waiting time in the queue under the two policies. This is similar to Figure 5. Figure 11 describes this information for messages that were not in the head of the queue (obviously relevant only with the exhaustive policy). Figure 10 displays the average time a message waited from the time it arrived to the time its last bit left the node. Clearly, this is similar to the times in Figures 5 and 9.

- The number of packets that were transmitted per node for each arrival of a token is also a metric of interest. Figure 12 describes this metric ($E(L]$). Not surprisingly, the patterns in this figure are similar to the ones in Figures 5,9, and 10. The plots for the non-exhaustive policy are not relevant since it only transmits at most one packet each time with a token.

- Network utilization is a very important metric in determining the ability of a protocol to function at high loads. Figure 13 summarizes this metric for the two policies. With the non-exhaustive policy, with a single ring the network (actually with a capacity of 1 gigabit/sec), can't have more than 60% utilization of the network capacity (even when the input load is high). This is certainly a restriction. Similarly, a two-ring network cannot have more than 75% utilization. The exhaustive policy, however, seems to place no such restriction, and the utilization appears to be independent of the number of rings.

4 Conclusions

From the above results, we make the following conclusions:

- At all loads, irrespective of the policy, there is a gain in having multiple rings. Whether this gain reduces when the number of rings increases (and hence the bandwidth of each ring decreases) beyond a certain point, is yet to be seen. We propose to experiment with 32, 64, and 128 rings to make stronger conclusions about this impact.

- Gigabit speed networks with a small number of parallel rings (1,2, or 3) will limit the utilization of the network. Thus, the total capacity of the network can never be utilized. The exhaustive policy, however, seems to perform well with any number of rings.

- The reduction in response time seems to be significant with multiple rings and at high loads. Certainly, there is a noticeable reduction in response even at medium loads with the multiple ring structures.

- The token inter-arrival seems to be significantly affected by the number of rings and the choice of the policy. This fact is very significant when parallel networks are used to support real-time applications.

In summary, the results are very interesting, but some more studies with other policies, bigger networks (more nodes, larger lengths), and more rings need to be carried out before a well established guidelines are set towards the design of parallel gigabit networks.

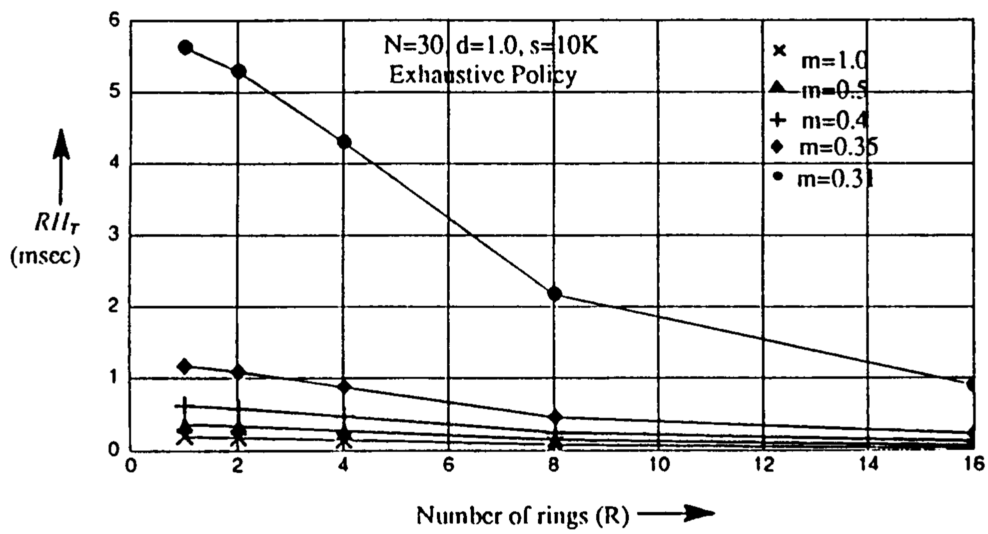

Figure 1. Mean Residual Time for Messages at the Head of the Queue

Figure 2a. Exhaustive Policy

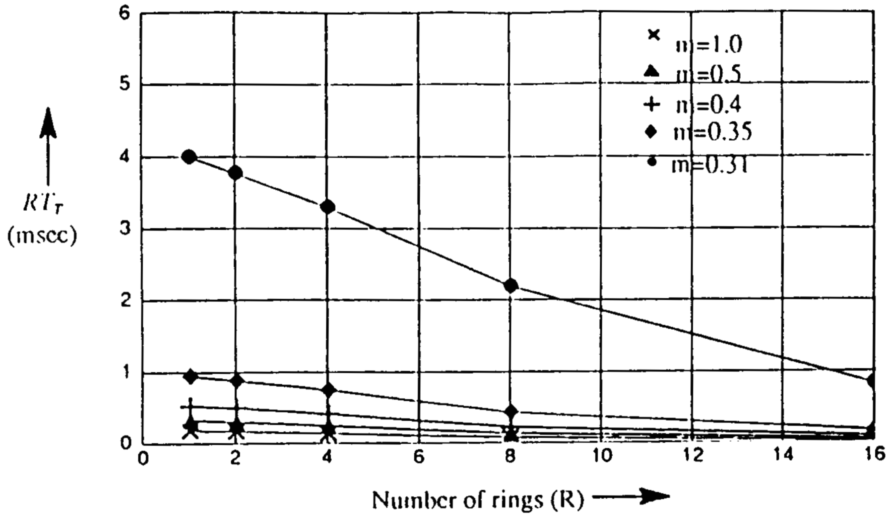

Figure 2b. Nonexhaustive Policy

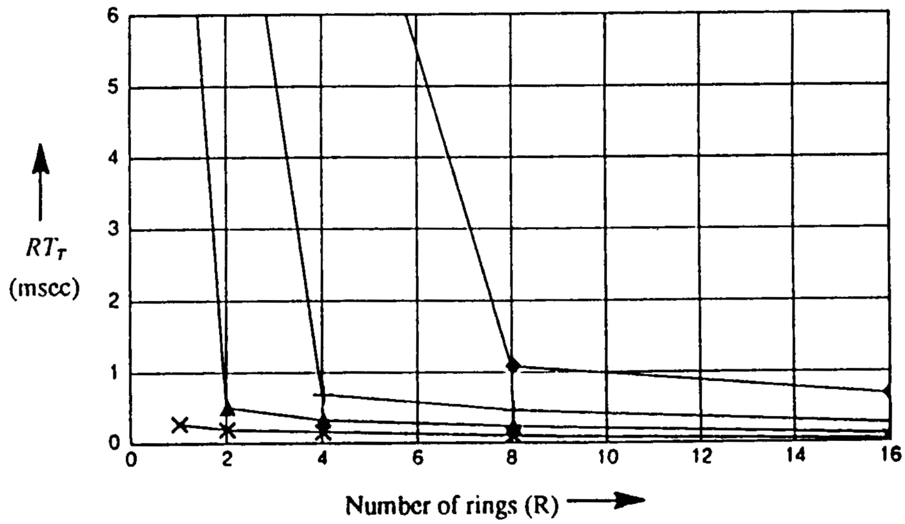

Figure 2. Average wait time (of a message) in a queue before a relevant token was received by its source node (N=30, d=1.0, s=10K)

Figure 3a. Exhaustive Policy

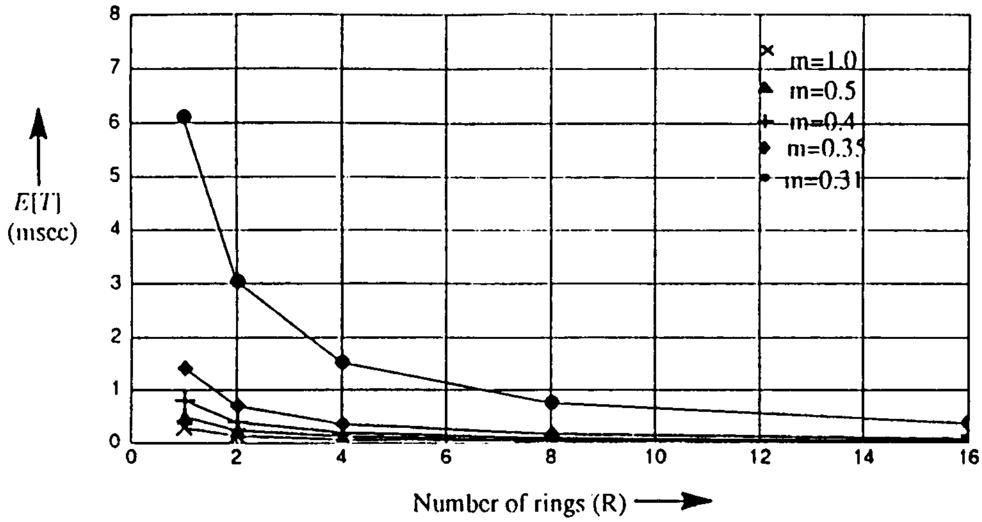

Figure 3b. Nonexhaustive Policy

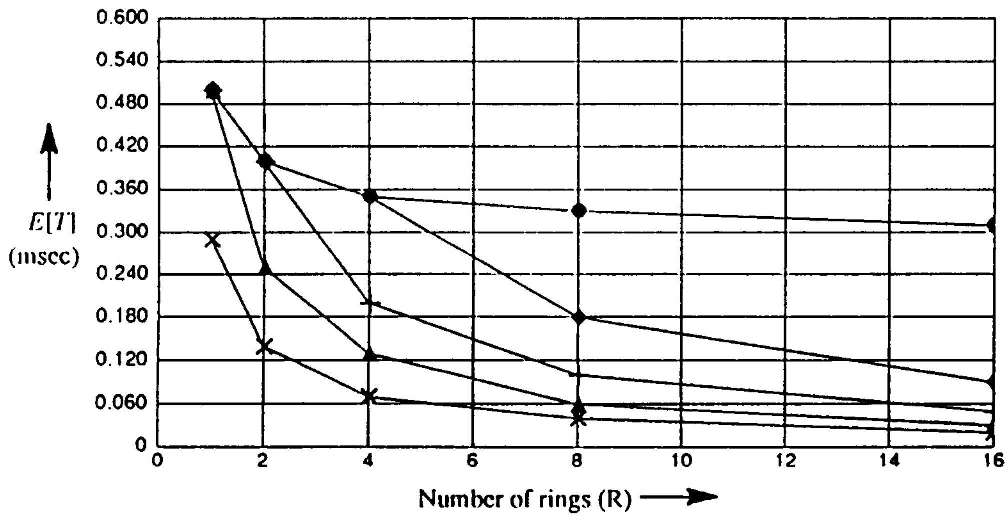

Figure 3. Average Token Inter-arrival Times at Each node (N=30, d=1.0, s=10K)

Figure 5a: Exhaustive Policy

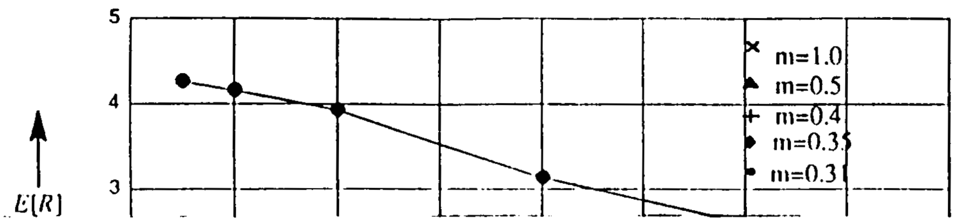

Figure 4a: Exhaustive Policy

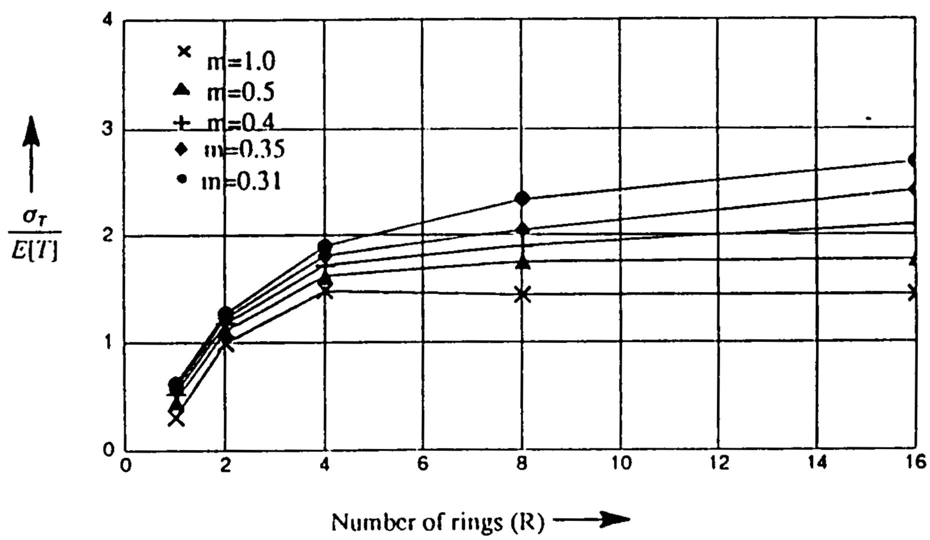

Figure 4b: Nonexhaustive Policy

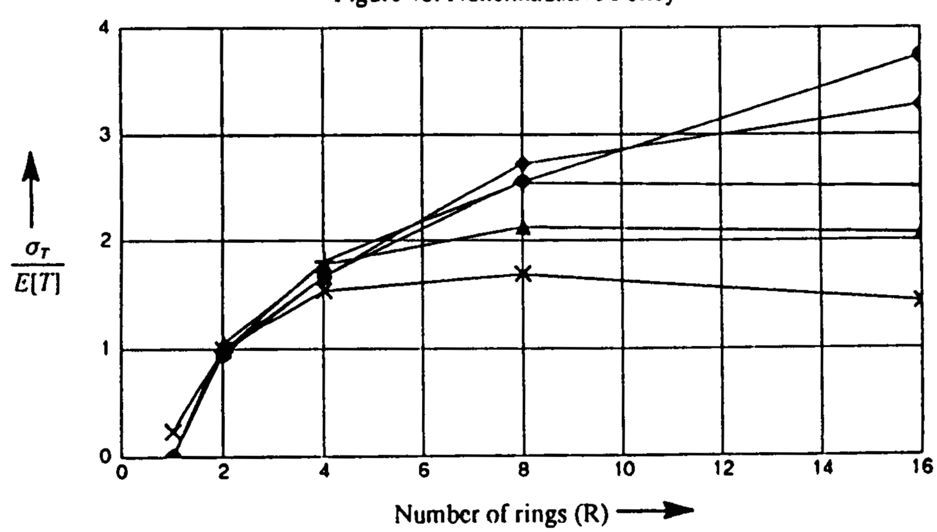

Figure 4. Standard deviation/Mean ratios for Token Inter-arrival Times

(N=30, d=1.0, s=10K)

Figure 6a: Exhaustive Policy

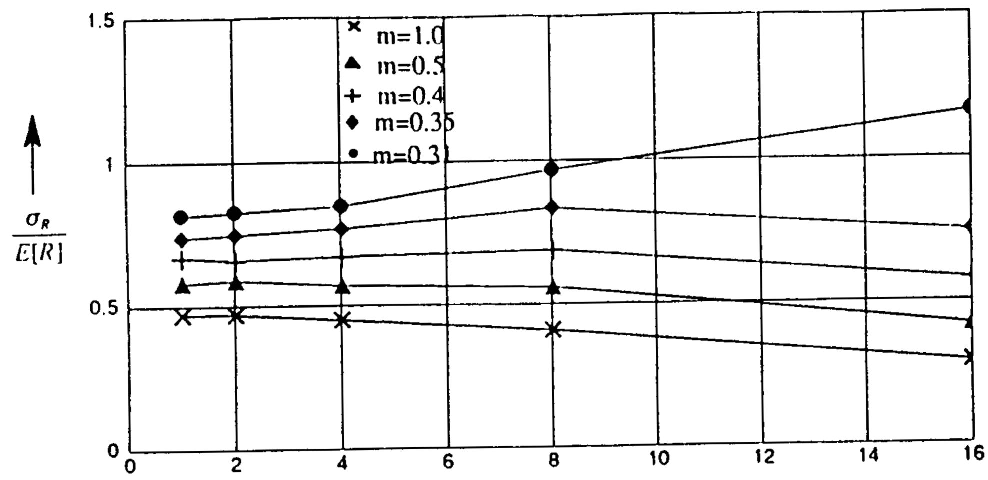

Figure 6a: Nonexhaustive Policy

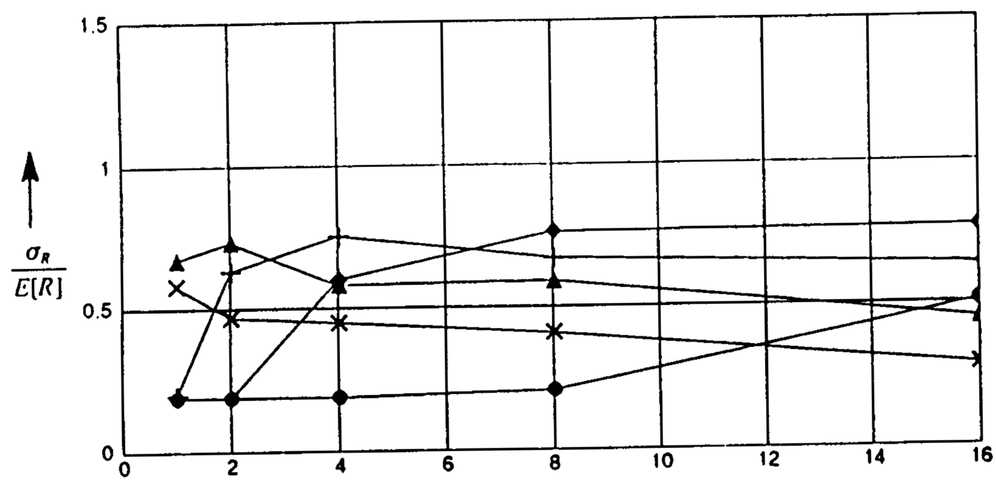

Figure 6. Standard deviation/Mean ratio for Packet Response Times (N=30, d=1.0, s=10K)

Figure 7a: Exhaustive Policy

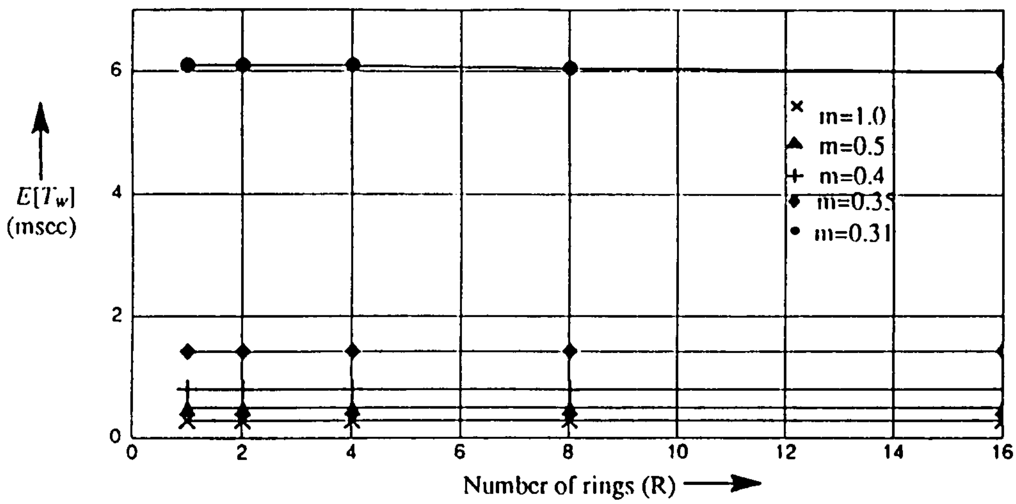

Figure 7b: Nonexhaustive Policy

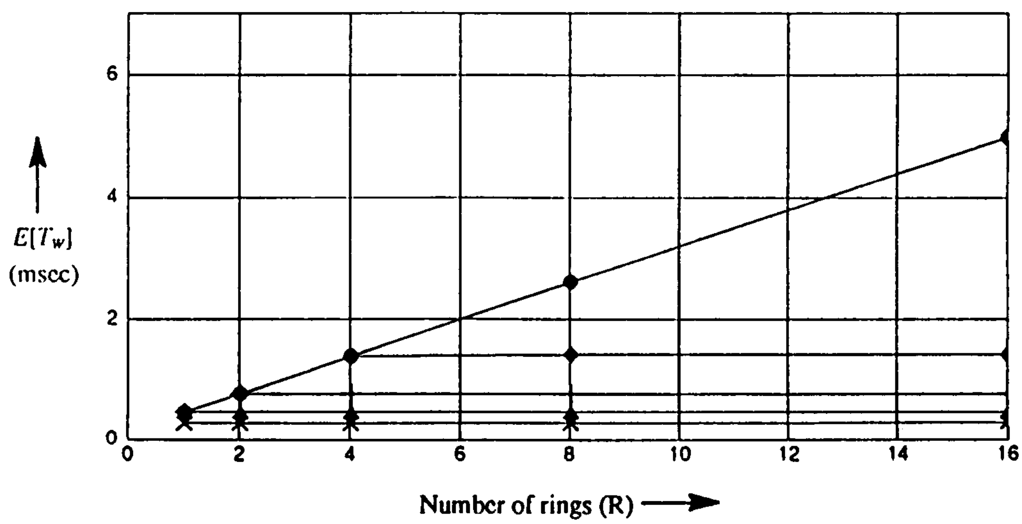

Figure 7. Average Token walk-times for each Ring (N=30, d=1.0, s=10K)

Figure 8a: Exhaustive Policy

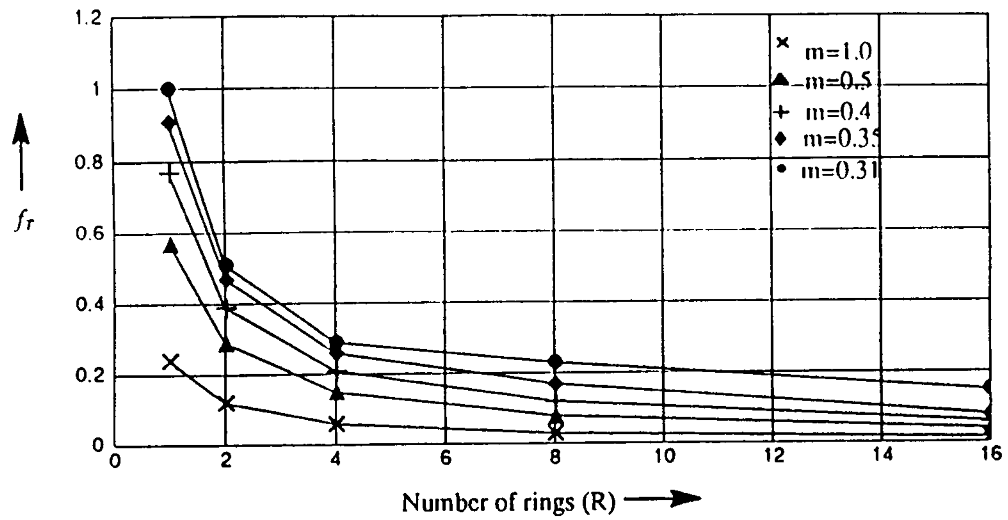

Figure 8a: Nonexhaustive Policy

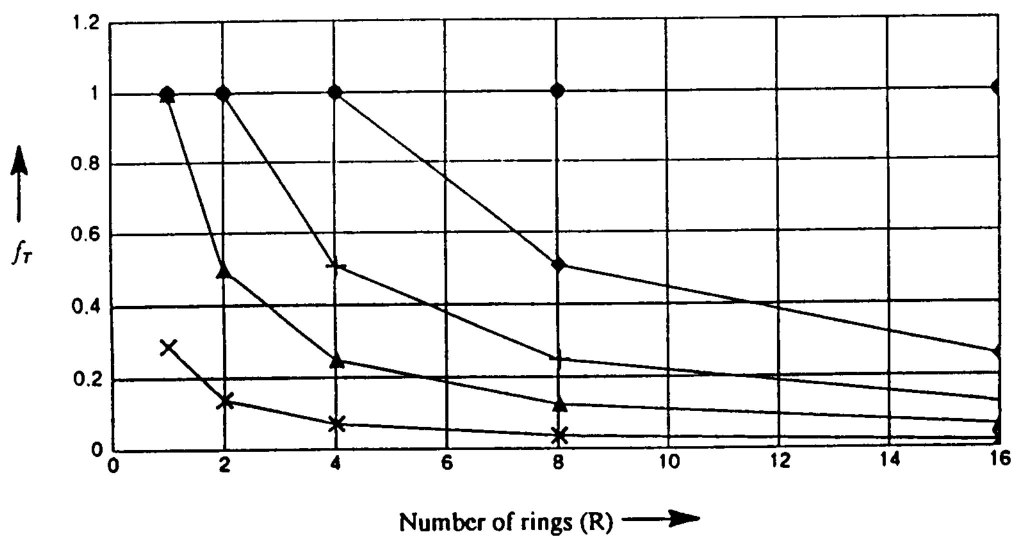

Figure 8. Fraction of times that an Arriving Token is Used by a Node for Transmission

(N=30, d=1.0, s=10K)

Figure 9a. Exhaustive Policy

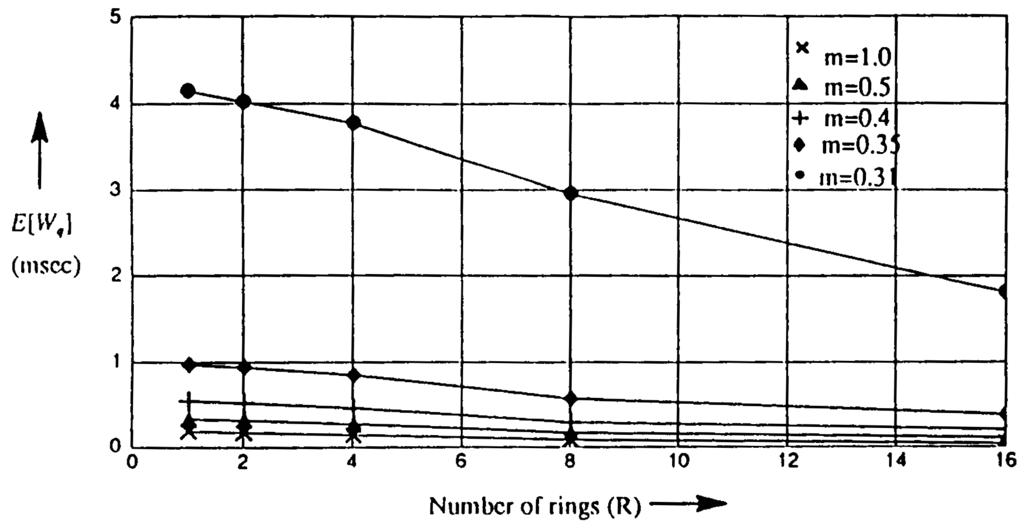

Figure 9b. Nonexhaustive Policy

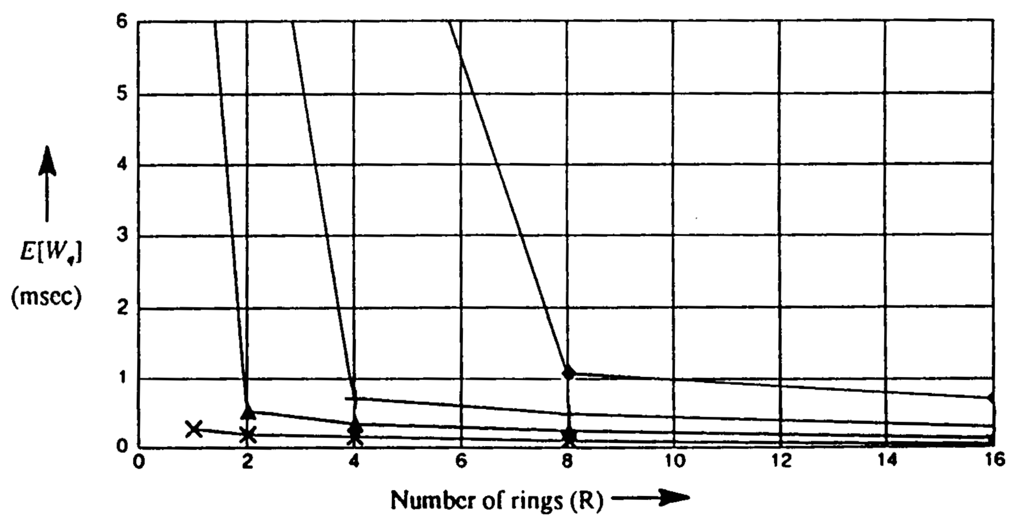

Figure 9. Average Waiting Times for Each Message (at a Node)
(N=30, d=1.0, s=10K)

Figure 10a. Exhaustive Policy

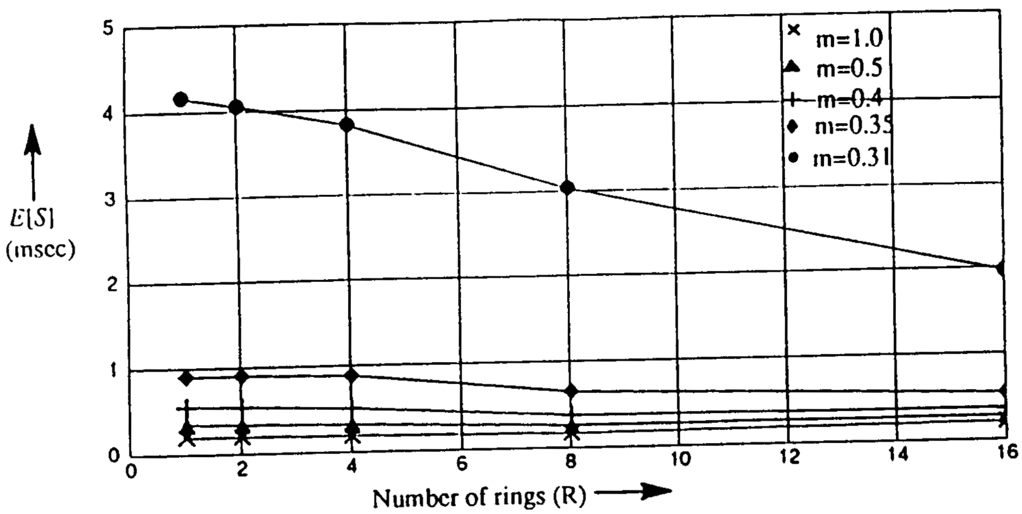

Figure 10b. Nonexhaustive Policy

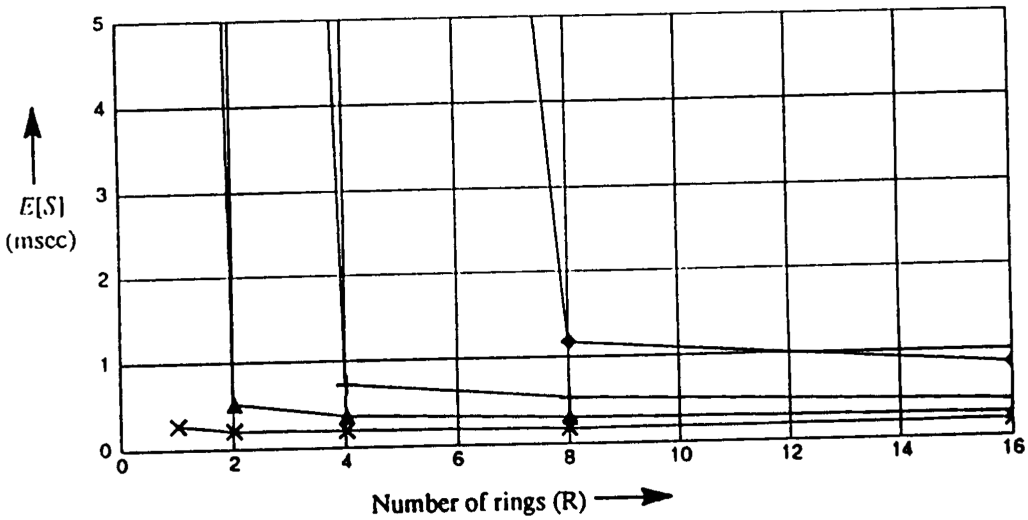

Figure 10. Average Time Between Message Arrival and Its Leaving the Node

(N=30, d=1.0, s=10K)

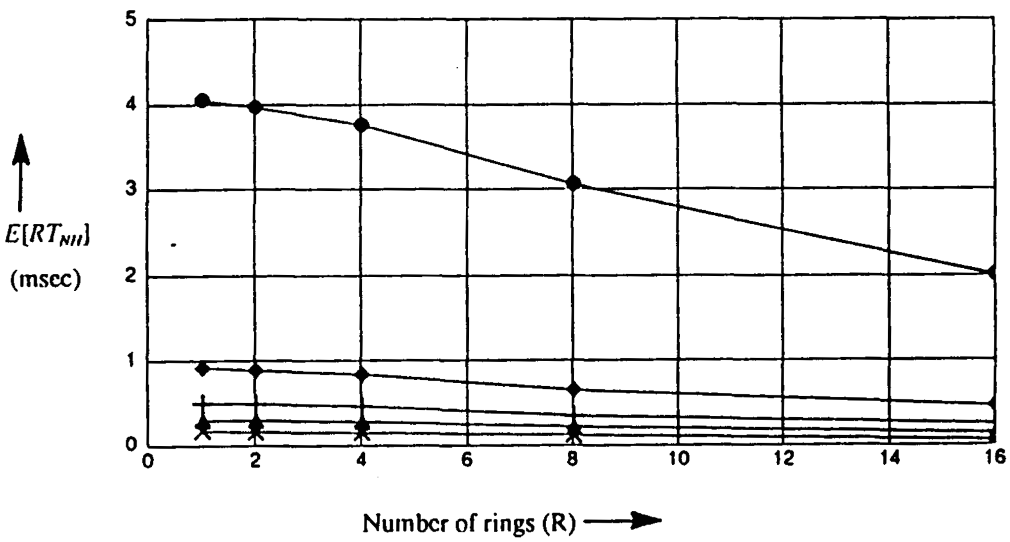

Figure 11. Average Response Time for Non-header Messages in a Queue
(Exhaustive Policy; N=30, d=1.0, s=10K)

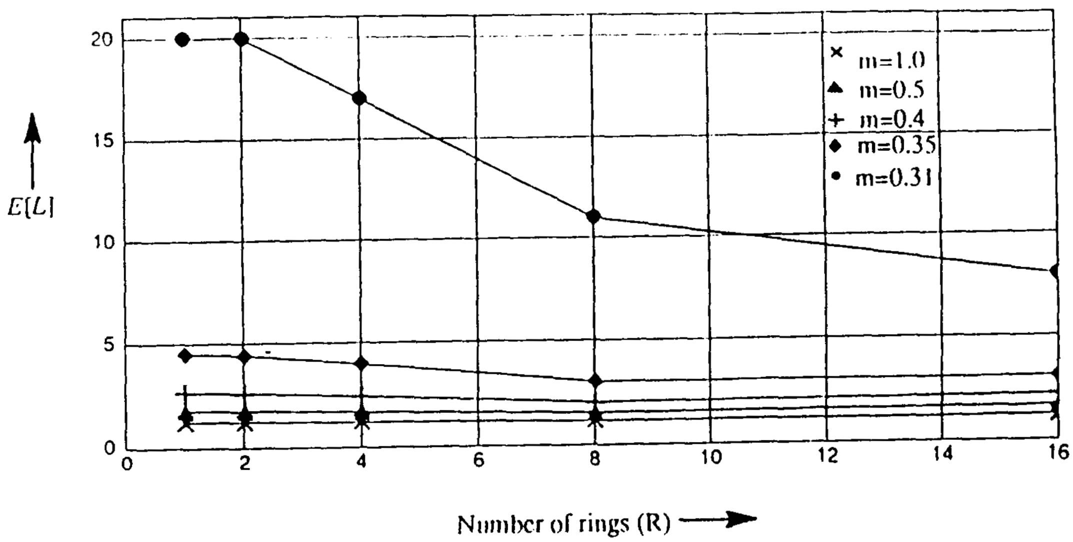

Figure 12. Average Number of Messages Transmitted per Node Per Token

(Exhaustive Policy; N=30, d=1.0, s=10K)

Figure 13a. Exhaustive Policy

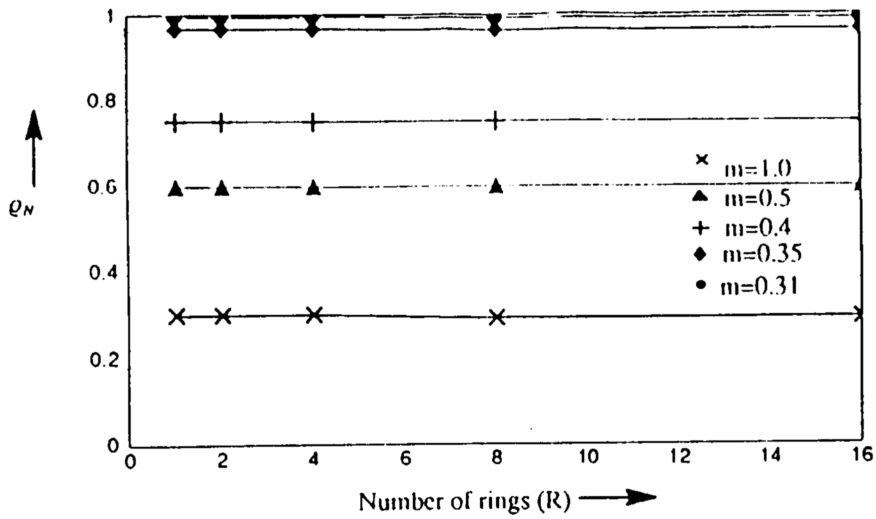

Figure 13b. Nonexhaustive Policy

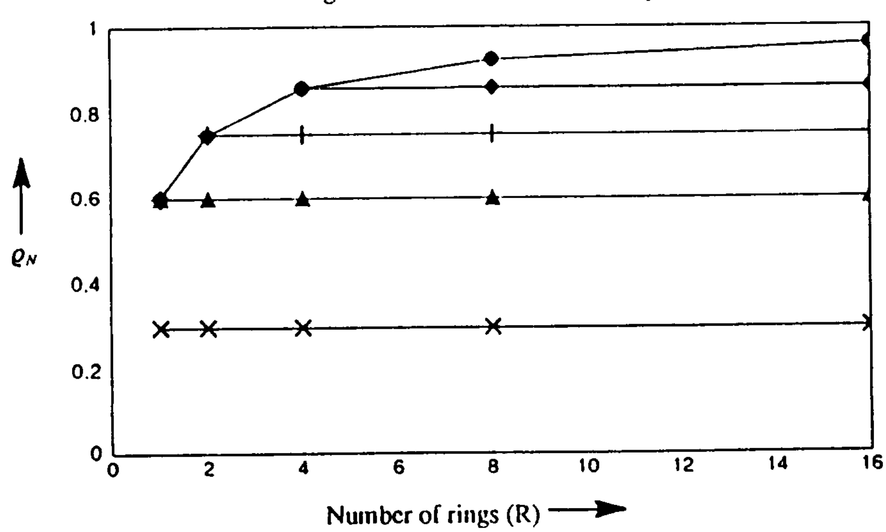

Figure 13. Network Utilization under the Two Policies

(N=30, d=1.0, s=10K)